## Robert J. Deri, Center Leader

The Center for Complex Distributed Systems is a focus for advanced technology development and applications involving ultrascale, spatially distributed systems. The Center seeks to develop an advanced suite of technical capabilities that enable the engineering of such systems, and to provide these capabilities in support of missions and programs at Lawrence Livermore National Laboratory (LLNL).

"Complex distributed systems" can be loosely defined as aggregations of large numbers of coupled, cooperating elements or subsystems, in which the system behavior cannot be described by simple hierarchical or nearest-neighbor interactions between elements. Examples include distributed sensor networks and beam diagnostics, large-scale distributed control and communication systems for optical and/or accelerator beams, and distributed information processing systems based on an underlying distributed sensor or control system.

Complex distributed systems play a critical role in several of LLNL's national security missions, with the potential to impact LLNL programs in other areas as well. This is illustrated by the four projects described in this report, all of which address applications of national importance and current interest:

1. "Accelerating the Development of Petaflop Applications and Systems" explores extremely large computational platforms, for application areas such as weapon simulations and climate modeling.
2. "Combined Sensing and Simulation for Enhanced Evaluation of Structures" describes the use of networked sensors to monitor the robustness of critical structural systems, such as bridges, dams, and buildings.
3. "Modeling and Simulation of Wireless Sensor Networks" investigates wireless sensor networks for applications such as early warning against attack by chemical agents. The networked sensor technology described in this work, and in the preceding article, will also impact environmental and industrial monitoring.
4. "Information Warfare Analysis Capability" discusses the challenges associated with, and techniques required for, defending information systems against attack.

A wide variety of technical disciplines are required for ultrascale system engineering. At this time, the Center is focussing on the areas of communications and control, system engineering, and simulation. These areas present fundamental challenges for most large-scale systems, as illustrated by all projects described in this report.

Communication and data flow issues are central to all four projects, as is the ability to simulate and design ultrascale systems. This underlying commonality of basic issues indicates that the technological capabilities developed by these investigators is generally applicable and transferable to other applications for complex distributed systems.

# Center for Complex Distributed Systems

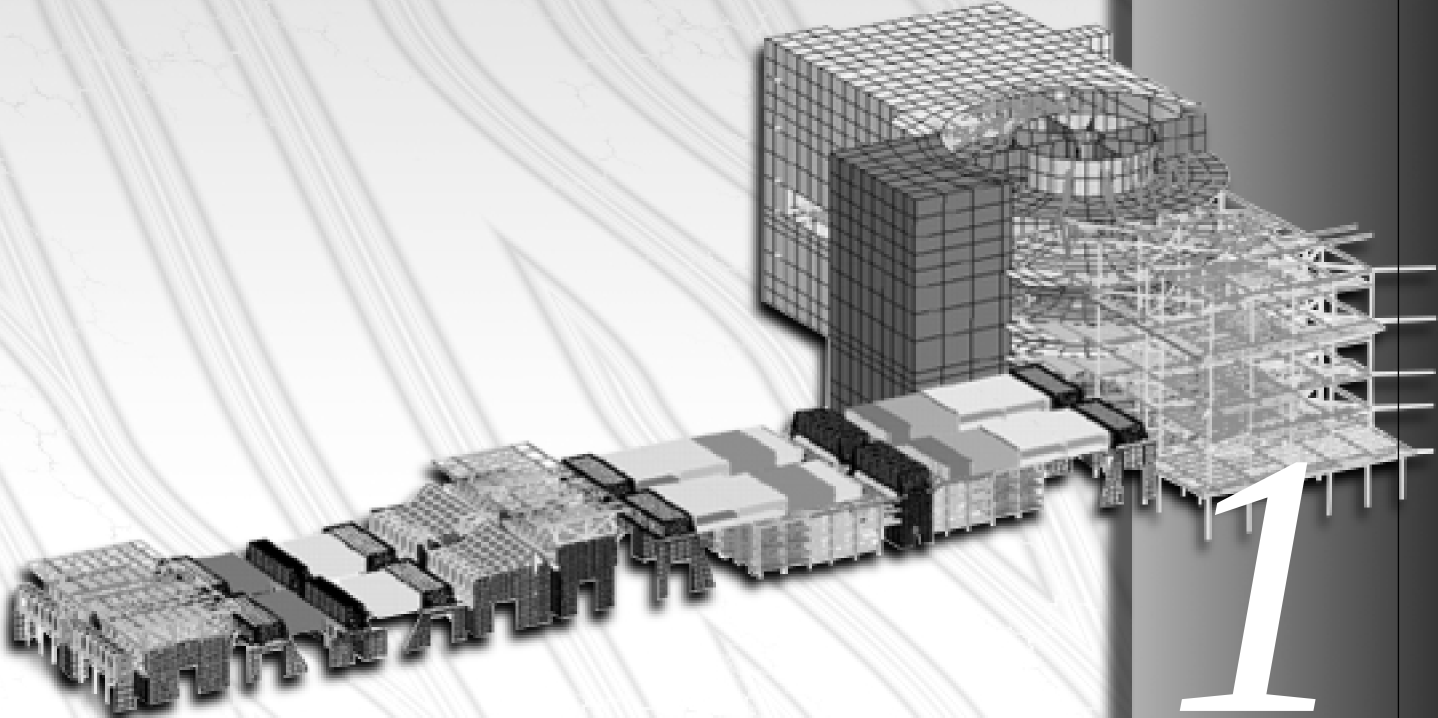

1

# 1. Center for Complex Distributed Systems

# Accelerating the Development of Petaflop Applications and Systems

**Anthony J. De Groot**
*Defense Sciences Engineering Division*
*Electronics Engineering*

**Robert J. Deri**
*Electronics Engineering Technologies Division*
*Electronics Engineering*

We are creating a modeling and analysis framework to accelerate the development of petaflop applications and systems.

## Introduction

The purpose of this project is to accelerate the development of technologies necessary to the creation of petaflop scientific applications and petaflop computing systems. This project brings together a multi-disciplinary team at Lawrence Livermore National Laboratory (LLNL) (computations, defense, engineering, physics, and lasers) to develop a modeling and analysis framework that will predict performance and identify bottlenecks for highly-relevant, unclassified Accelerated Strategic Computing Initiative (ASCI) applications (scaled to petaflop-class) running on proposed commodity-based architecture designs.

This will give LLNL a new capability that could assist in developing new high-performance applications, evaluating new hardware acquisitions, and creating a balanced computing environment.

## Progress

One of the activities for this fiscal year was to begin development of a simulation capability to understand how interprocessor communication affects performance.

Simulation of the execution of large, complex applications on petaflop computing systems is a computationally intensive task. In the past, we have simulated the execution of an application at the instruction level. The computation required for this highly detailed simulation is approximately one hundred times more than the computation time to run the application on a real machine. We chose trace-based simulation to make this problem computationally tractable.

In our trace-based simulation, the application with its input problem are executed on a host machine, producing a trace of events. Events include all message passing operations, and optional user-defined events, including key subroutine entrances and exits. This trace, along with a description of the computer architecture to be simulated is used as input to a discrete-event simulator, that simulates the execution of the application on the simulated architecture.

Generation and visualization of traces and subsequent analysis have already improved performance of a complex application. We chose Parallel DYNA3D (ParaDyn) as the first application to study because it is one of the few parallel production codes far enough along in development for this project to make an immediate impact. We used the AIMS code from NASA Ames to generate and display the event traces from ParaDyn execution.

**Figure 1** shows the execution of ParaDyn on eight processors. The problem simulated by ParaDyn is that of a buckling beam. In the figure, the horizontal direction represents the passage of time, and the vertical direction represents processor number. Line segments connecting bars represent message communication from the corresponding processors. The shading of a bar represents the execution of a set of subroutines by the corresponding processor. **Figure 1** shows that most of the time is used for computation and very little is used for

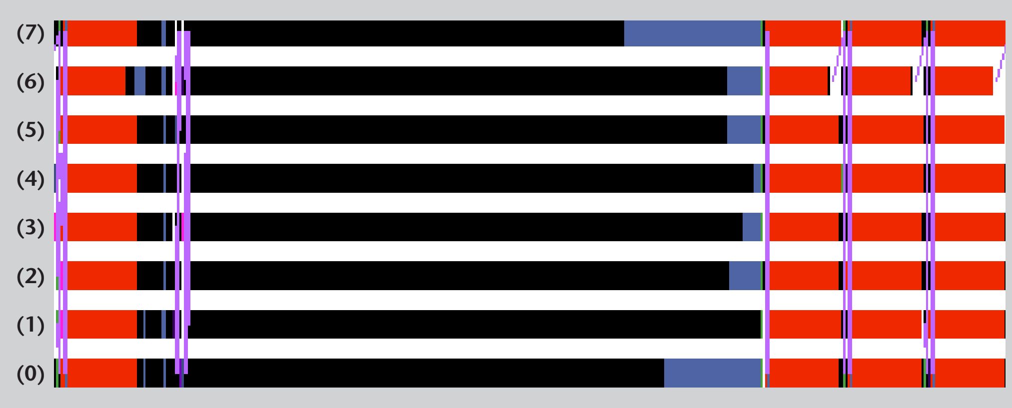


*Figure 1. Visualization of ParaDyn execution trace, showing excessive time spent in the subroutine, represented by the black bars.*

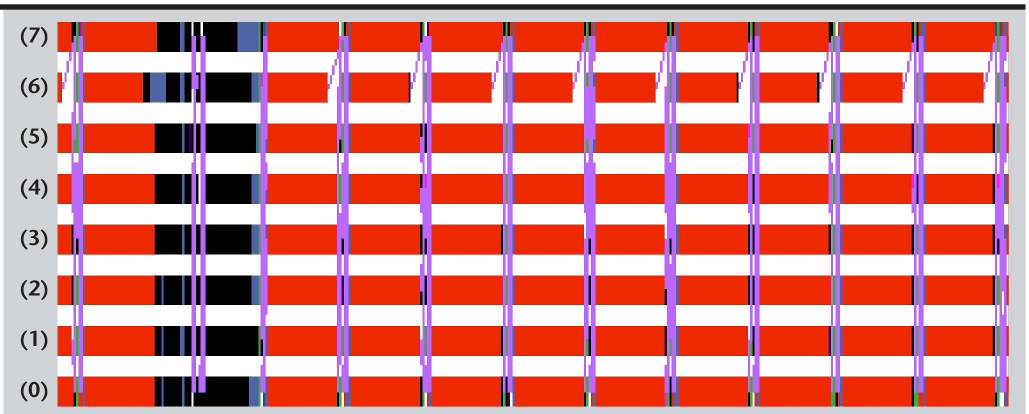


*Figure 2. Visualization of ParaDyn execution trace after program optimization suggested by Figure 1.*

communication. The figure also shows that the execution of the subroutine represented by the black bars is excessively long.

Further examination of the code showed that an optimization could be made to that subroutine to reduce execution time. **Figure 2** shows the result of that optimization, where the total execution time of that ParaDyn problem was reduced by a factor of about 1.5. Further optimizations may be possible with additional study.

# Combined Sensing and Simulation for Enhanced Evaluation of Structures

**David B. McCallen and Matthew Hoehler**
*New Technologies Engineering Division*
*Mechanical Engineering*

**Gregory A. Clark and James V. Candy**
*Electronics Engineering Technologies Division*
*Electronics Engineering*

Large-scale computer simulation is an essential tool in the design and analysis of modern structures. With the enormous cost and construction effort associated with many large structures, it is imperative that computer simulations provide an accurate picture of structural performance spanning a multiplicity of loading environments such as bomb blasts, earthquakes, and ambient vibrations. In the current study, techniques are investigated that allow evaluation of simulation model accuracy and the possibility for subsequent enhancement. The multidisciplinary tools being used at Lawrence Livermore National Laboratory (LLNL) include finite-element-based structural simulation, model-based signal processing, and remote sensing and data communication. The overall objective of this research is to symbiotically couple numerical simulation with field measurement of structural behavior. The result will be enhanced accuracy and reliability of numerical simulation of structural response, and the ability to monitor fundamental changes in complex structural systems, a prerequisite to health monitoring and damage detection.

## Introduction

With the proliferation and advancement of sophisticated numerical simulation software tools over the past twenty years, computational simulation of large structural systems has been a subject area experiencing rapid growth. Engineers now rely heavily on large-scale structural simulations to design and evaluate the performance of critical new structures, and to establish the performance of expensive retrofits on existing structures.

Despite the advances in computational methods, there remains a significant degree of uncertainty in predicting the field performance of many large-scale structural systems. These uncertainties are rooted in our inability to precisely quantify the phenomenological behavior of certain aspects of structural excitation and structural response—for example, uncertainties in the precise deformation characteristics of complex structural element interconnections or uncertainties in estimation of the actual loads a structure will be subjected to. To advance our ability to accurately and confidently simulate the response of structures, we must make use of measured structural response characteristics and measured excitation functions.

The notion of coupling simulation and measurement is emerging in the study and characterization of many complex systems found in the sciences and engineering. One of the tenets of a new National Science Foundation program on Knowledge and Distributed Intelligence is the recognition that: "Better understanding of complex phenomena now requires interplay between computations and data, often in real time." The work described in this report provides the foundation at LLNL for establishing an interplay between computations and measured data for the specific case of large structural systems.

The essential links between simulation and measurement are shown in **Fig. 1**. As indicated, information from a numerical structural model and data from field measurements of structural response are fed into a model-based signal processor. The signal processing toolbox is used to evaluate whether the model and as-built structure are in agreement, or if there is a discrepancy between the dynamical behavior of the model and that of the actual structure.

In addition, research on the signal processing toolbox will attempt to discriminate the source of existing discrepancies between the simulation model and the as-built structure. Once discrepancies can be identified, the basis for damage detection in a structural system has been established. Experimental measurements before and after an event can be used to identify changes in the structural system, and identification of the source of the differences with a damage detector can be performed to investigate where damage has occurred in a large distributed structure.

## Progress

In the first year of work, we have assembled a multidisciplinary team, the combined expertise of which spans the required technologies in mechanical and electronics engineering. Our accomplishments include the following:

1. An understanding of the relationship between the structural simulation models of mechanical engineering and the state space models of electronics engineering has been established for large structures.
2. The protocols and requirements for information exchange between second order equations of motion of a computational structural model, and first order equations of motion of state space, have been developed.
3. The software for a model-based signal processing state space model module has been completed for performance of structural simulations, and a state space filter has been developed for detection of differences between the simulation model and a measured structural response.
4. In the area of sensing and communication, a new data acquisition system has been developed which will allow on-demand wireless communication of data via cellular phone from an instrumented structure.

*Figure 1. Model-based signal processing for linkage between numerical models and measured response.*

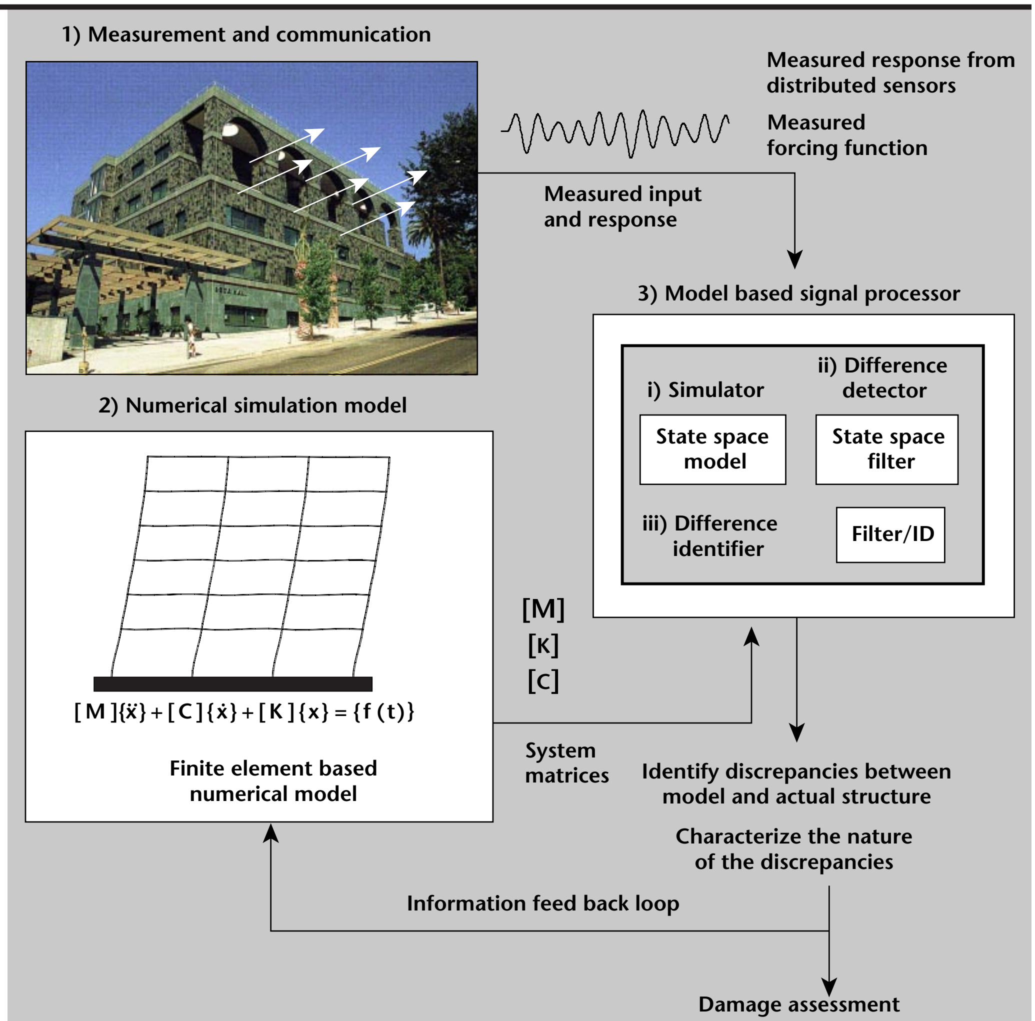

5. A field experiment has been completed at the Nevada Test Site to evaluate the data acquisition system and to provide a real world data set for evaluation of the model-based signal processing algorithms in year two of this project.

## Interfacing Numerical Simulation Models and Signal Processing

A capability for the model-based signal processing algorithms to detect discrepancies between a numerical structural model and the actual structure has been developed and coded in a MATLAB environment (**Fig. 2**).[1] The matrices constructed for the numerical simulation model are passed into the MATLAB-based state space model for process simulation (essentially a recast of the "N" second order equations of motion into a set of "2N" first order equations). Using the constructed state space model and measured data from a structure, a difference detector based on the whiteness test described by Candy[2] has been coded. This whiteness test, based on an autocorrelation check between measured structural response and the state space model,

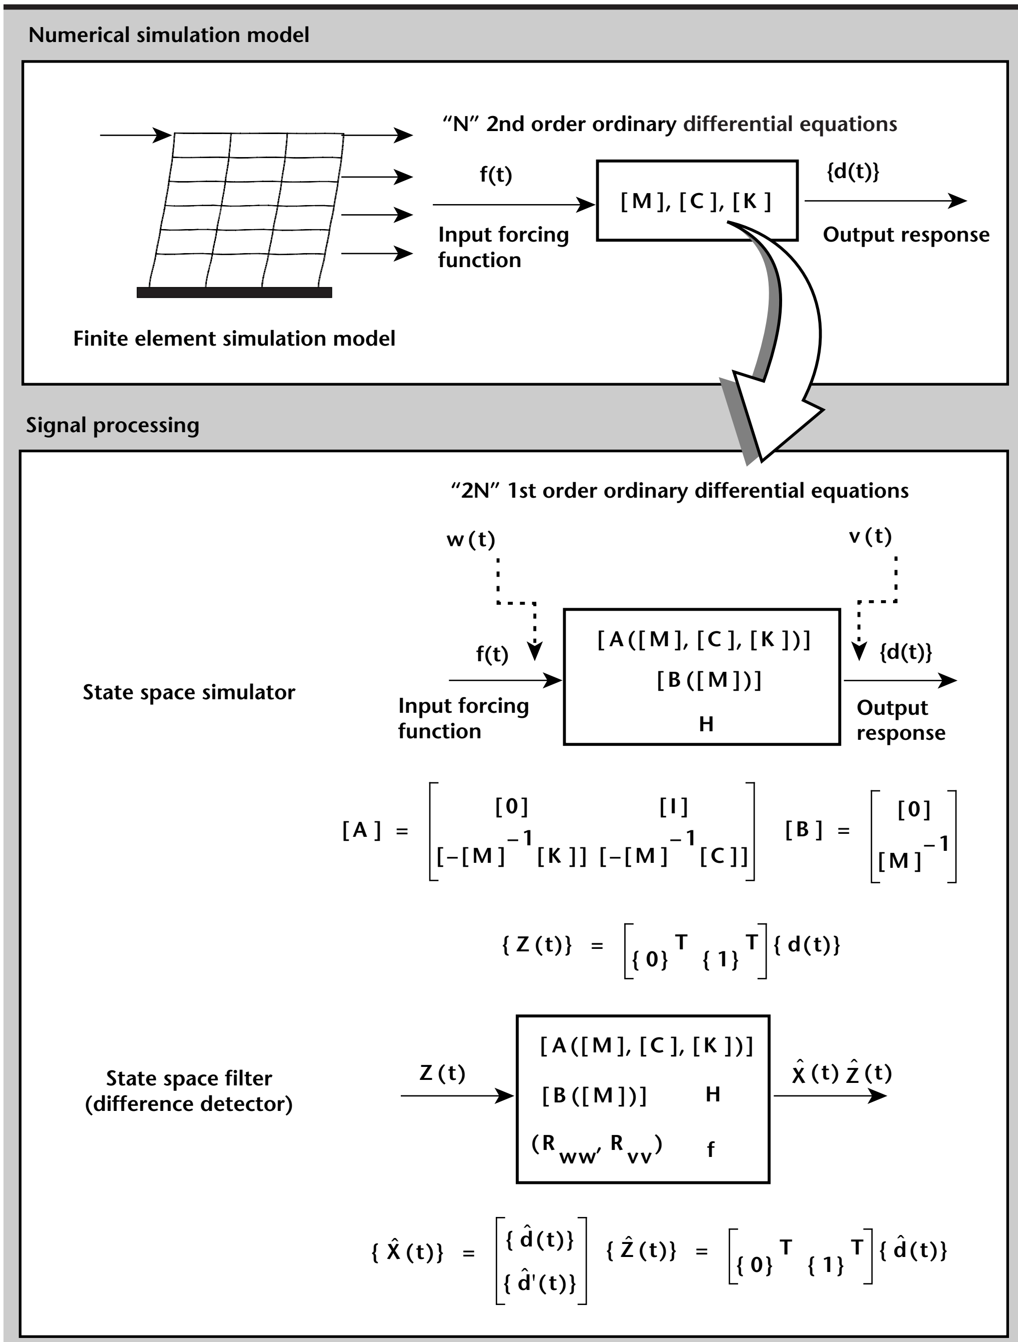


*Figure 2. Numerical simulation model matrix hand-off to signal processing state space simulator and state space filter.*

provides a rapid assessment of the validity of the numerical structural model.

An important feature of the model-based signal processor under development is the ability to work with full transient time history waveforms of the structural response. This is in contrast to most established model evaluation algorithms which are based on natural modeshape characteristics of the structure.[3] The ability to work directly with transient time histories will provide a powerful tool which can use any forcing function on the structure (that is, an impulsive load or sinusoidal load imposed on the structure by design). Furthermore, the signal processing methodology is not restricted to linear systems.

The ability of the signal processing algorithms to detect differences between the numerical structural model and the actual structure has been demonstrated with example problems. **Figure 3** illustrates an example for a five-story building, in which a simple five-degree-of-freedom "shear building" numerical simulation model of the structure has been constructed. In addition, a second model is constructed to be representative of the data from an actual "measured" structure. The second model is intended to provide the data which typically would be measured in the field.

To evaluate the ability of the signal processing algorithms to detect the difference between the model and the "measured" structure, transient response waveforms were determined for unit forcing functions applied to each floor level of the sample building. When the simulation model and "measured" structure are precisely the same structure, the whiteness test which discriminates differences between the model and actual structure is passed, indicating the model is in agreement with the actual structure (**Fig. 3a**). However, when the model representing the measured structure is perturbed by cutting the stiffness of the first floor level by one-half (**Fig. 3b**), the signal processing package immediately senses the difference between the numerical simulation model and the "measured" structure, providing the engineer with an immediate indicator of the deficiencies of the numerical structural model.

## Sensing, Communication, and Field Experimentation

The signal processing linkage between the numerical structural model and the measured structure requires the ability to economically monitor the response of a large structural system. For many large, distributed structures this can be a difficult, and potentially prohibitive task if traditional wired sensor systems are used. In cooperation with the private firm of Jarpe Data Systems, a data acquisition system has been developed that will allow remote gathering of transient response data from a large structure. The capabilities of this data acquisition system have been tested in an experiment at the Nevada Test Site (**Fig. 4**). The data acquisition system records a data stream from a suite of structural sensors, typically accelerometers, and stores the time-stamped data on a disk storage system. The acquisition system has a cellular phone system on board and can be called up remotely with a laptop computer with a modem for on-demand download of data. This data acquisition system will provide an economical and practical means of remotely monitoring large distributed structures.

## Future Work

The development of the model-based signal processing toolbox will continue with the construction of the Difference Identifier (the Filter/ID in **Fig. 1**). This will complete the basic signal processing package. Extensive sensitivity studies will be performed to assess the ability of the signal processing tools to detect and identify many different types of discrepancies between simulation models and measured structures. These studies will make use of computer simulation models (more sophisticated samples of the type shown in **Fig. 3**), existing full-scale structural tests performed at various universities, and the data obtained from the carefully controlled structural experiment performed at the Nevada Test Site in FY-98. In addition to identification of discrepancies, the research will determine the degree to which the precise cause of the discrepancies can be identified by the Difference Detector.

In parallel to the algorithmic developments, sensors and data acquisition systems will be placed on three large distributed structures to establish the ability to perform on-demand remote sensing over an extended time period. The target structures include the San Francisco-Oakland Bay Bridge[4,5] the National Ignition Facility laser bay structures, and the Bixby Creek Arch Bridge at Big Sur (**Fig. 5**).[6]

Each of these structures has been studied at LLNL as part of ongoing research and development projects or programmatic work, and numerical simulation models exist for each structure. These important structures present different sensing and communication challenges

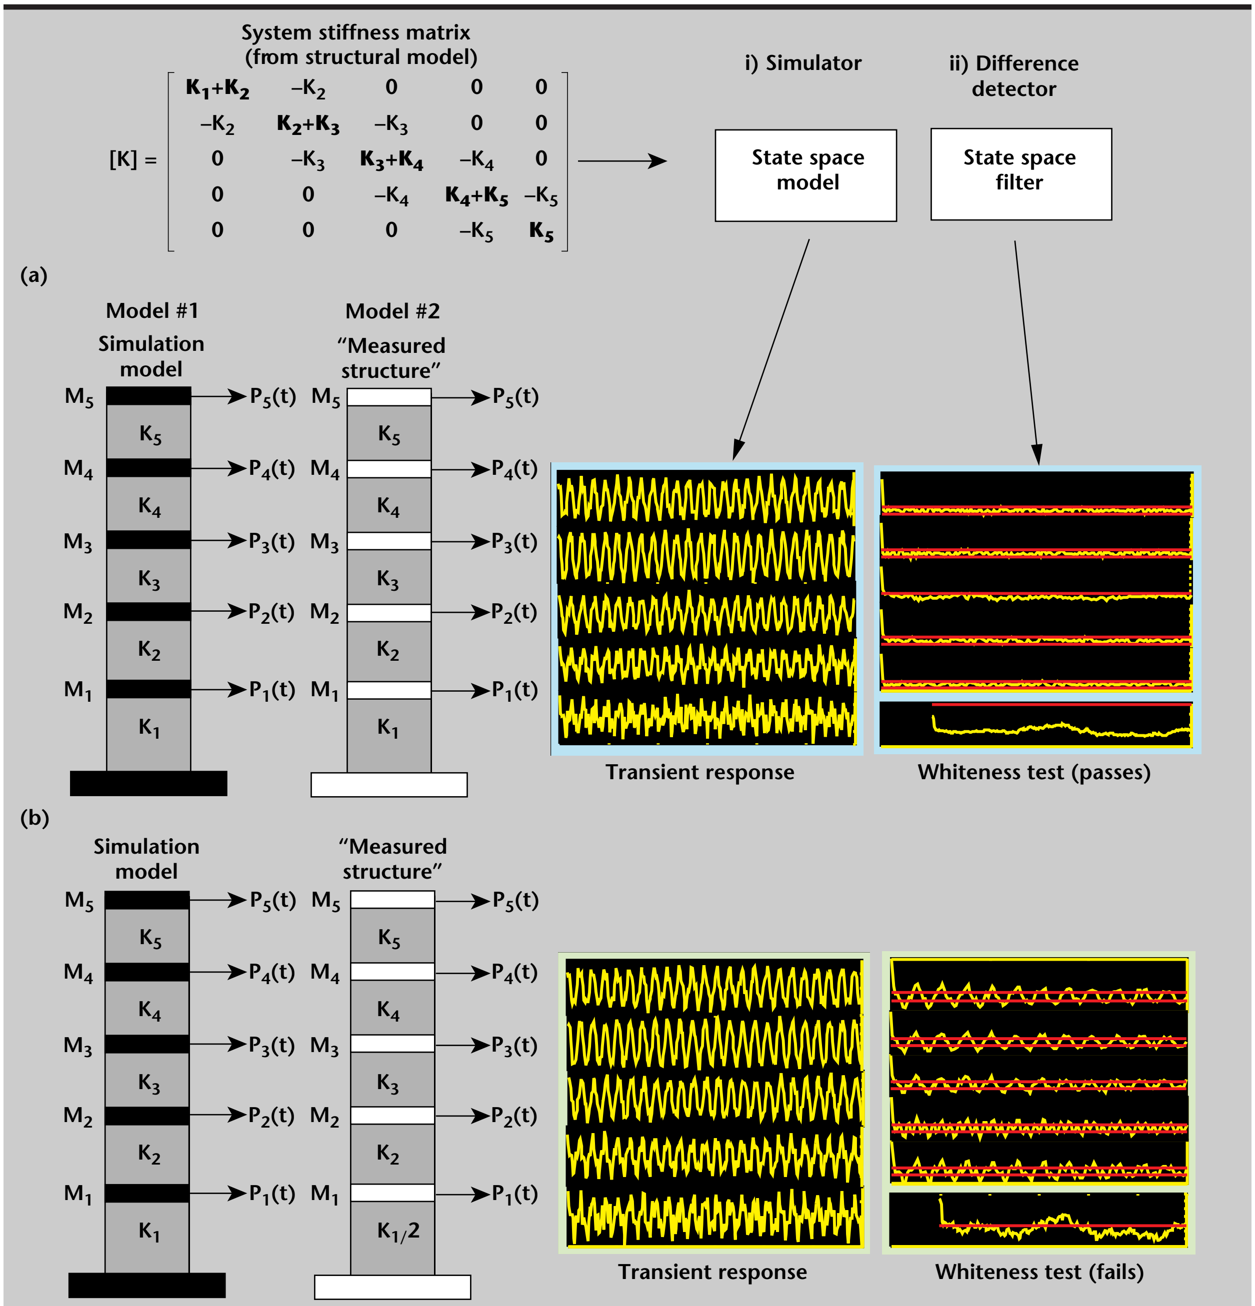


*Figure 3. Development of model-based signal processing components—a model-based simulator for process simulation and a state space filter for difference detection. a) Application to a problem in which the computational model and measured structure agree (whiteness test is passed); b) application to a problem in which the computational model and measured structure disagree (whiteness test fails).*

and will require structure-specific detailing of the sensors and data acquisition system. Previous study of these structures has provided the understanding of the frequency range of interest, and existing simulation models will be used to establish sensor placement locations. Successful monitoring of these structures will provide the hardware and methodology for deployment of dense instrumentation arrays on large distributed structures.

Large structural systems such as the National Ignition Facility or the Bay Bridge represent a tremendous capital investment and are critical structures for our society. The degree to which we analyze, design and monitor these structures should be commensurate with the cost and criticality of function which

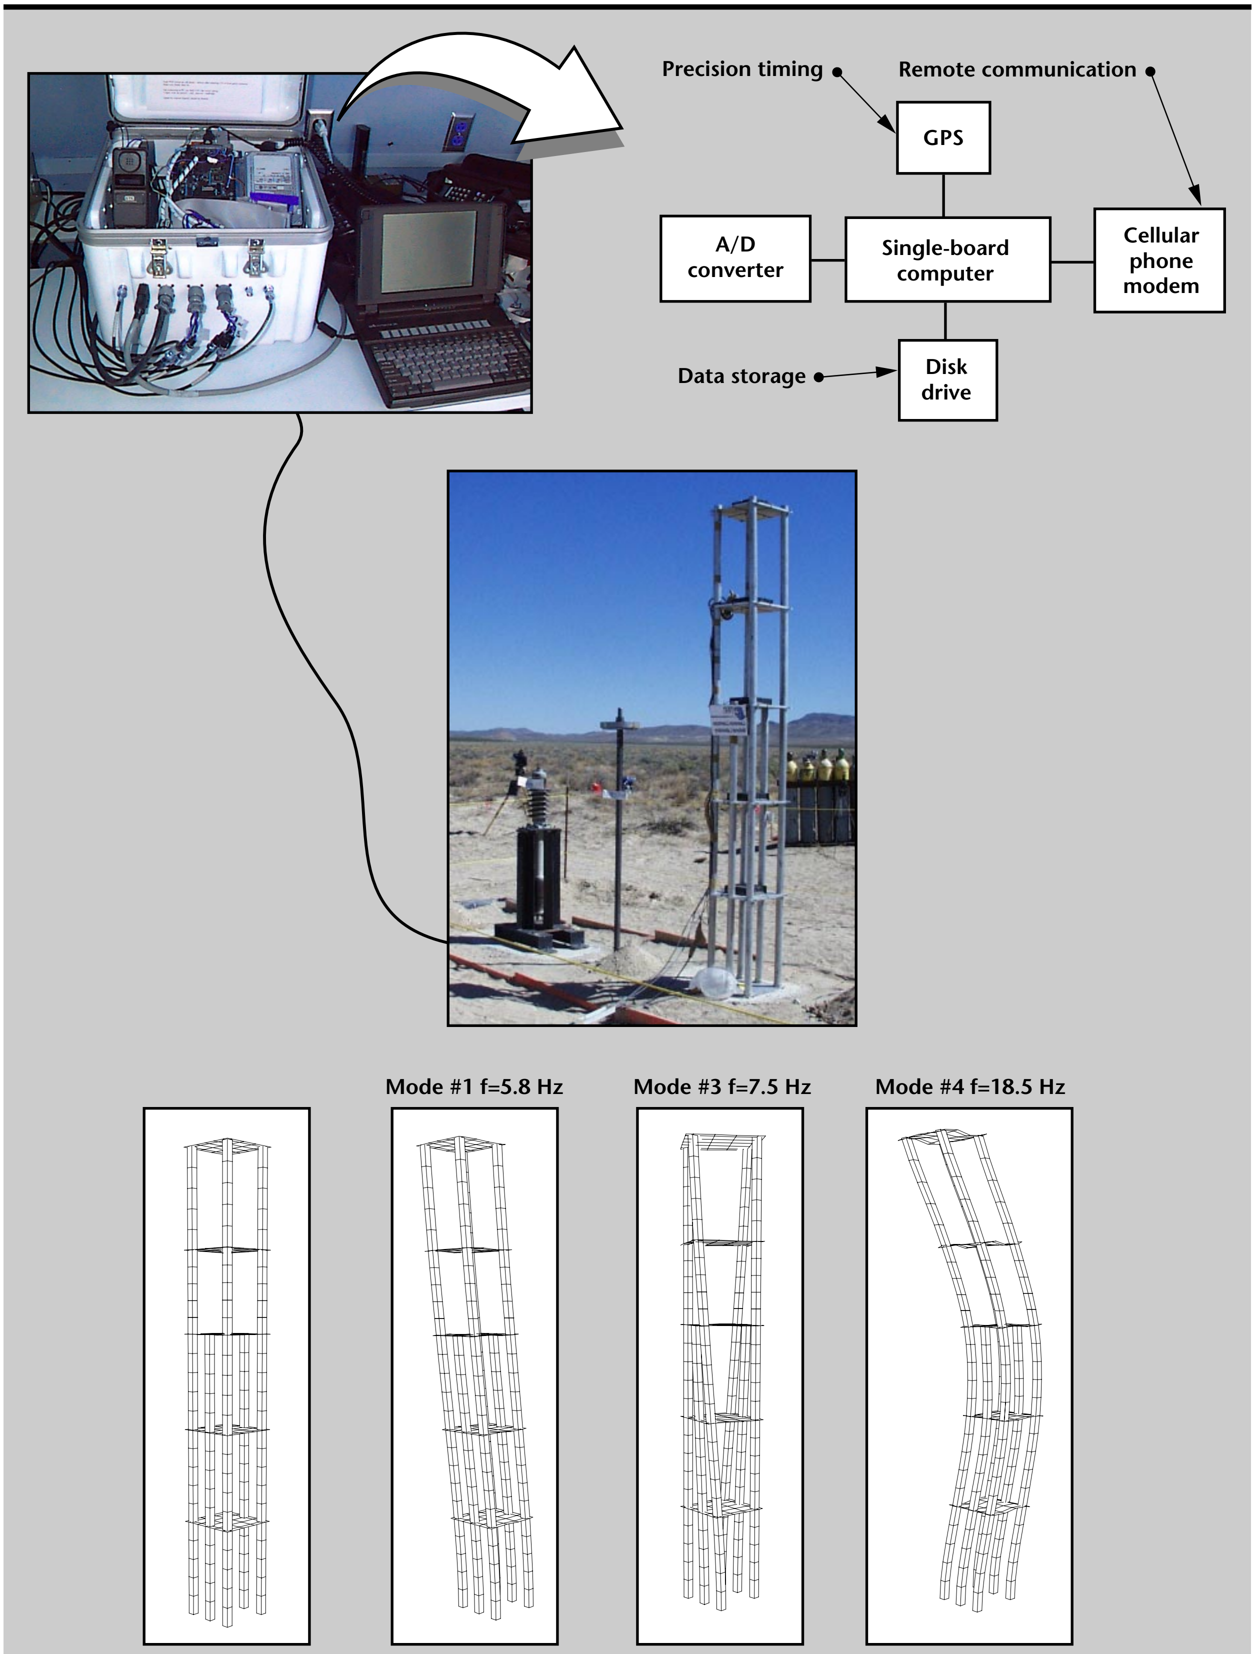


*Figure 4. Data acquisition system for remote, wireless monitoring of a distributed structure. System includes GPS for precision time-stamping of data, and on-board cellular phone communications.*

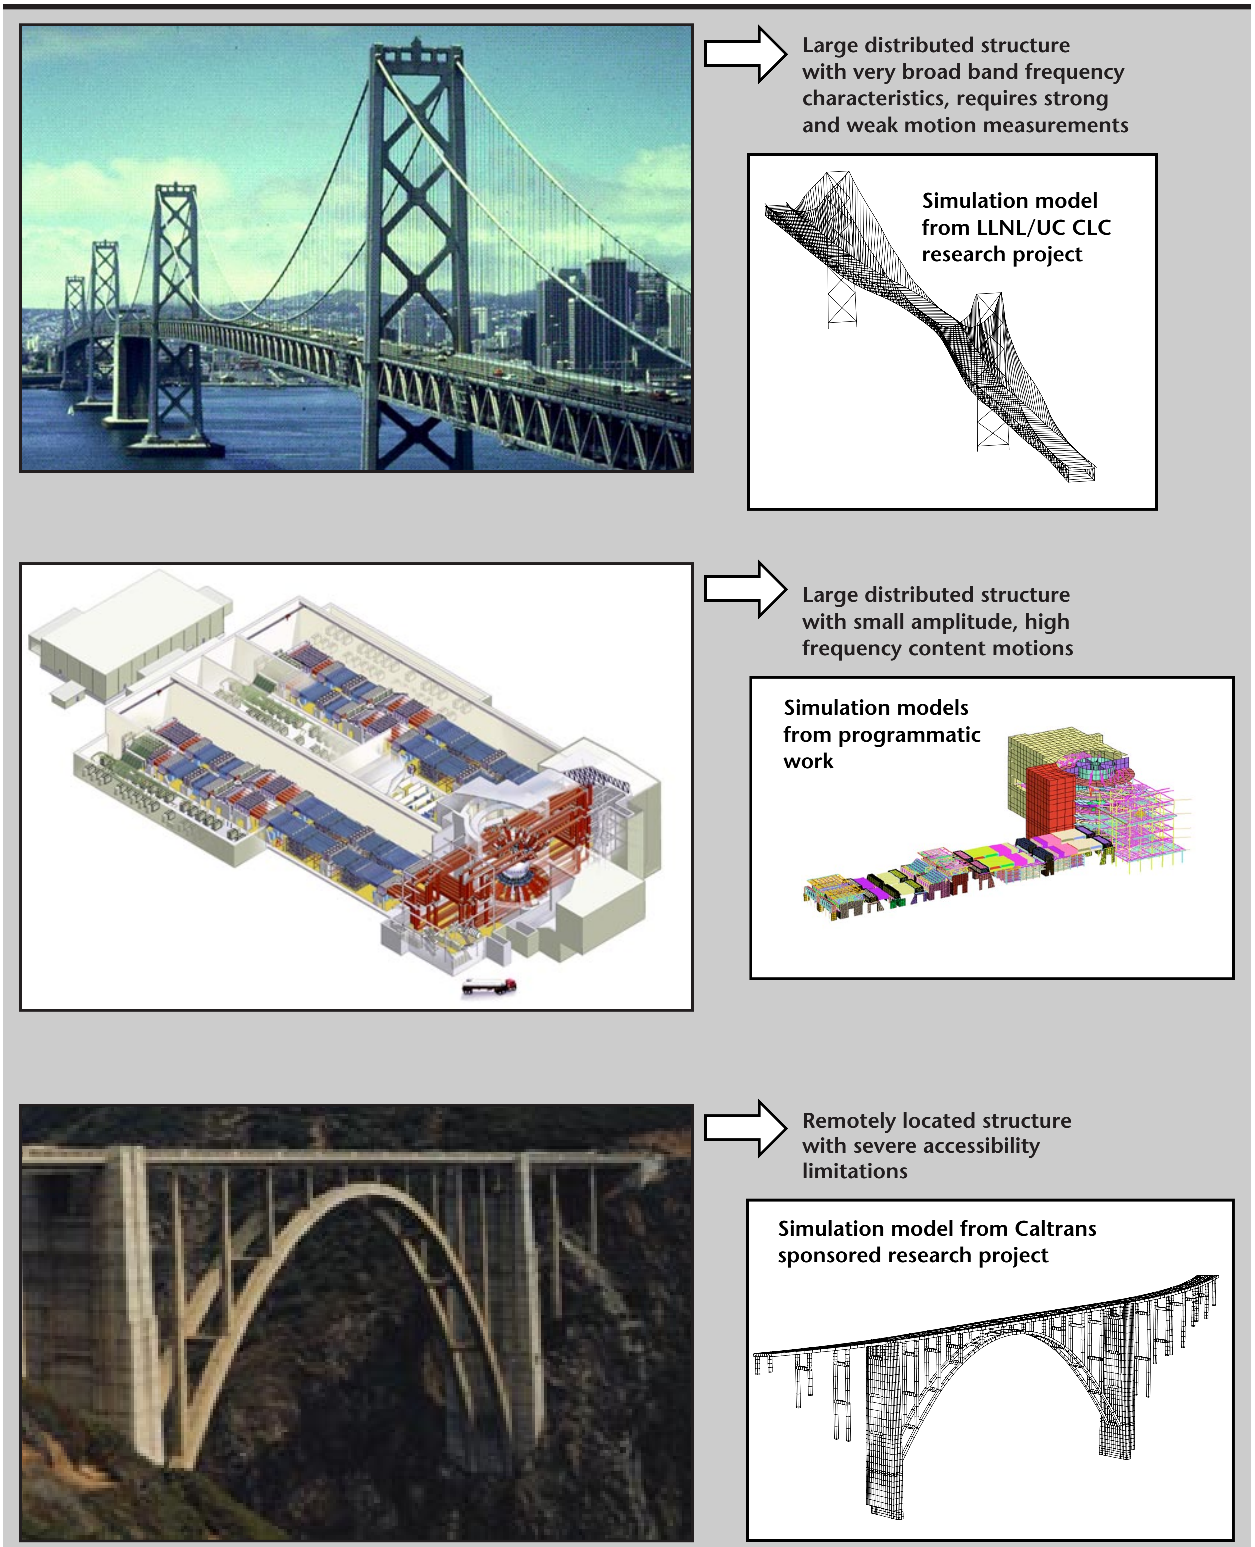


*Figure 5. Large structural systems to be instrumented in FY-99.*

they perform. The methodologies developed with this research will provide the tools necessary to significantly improve our understanding of how large distributed structures behave, will lead to increased reliability in numerical simulations, and will provide an entirely new capability for long term monitoring to ensure the integrity of important structures.

## References


1. Clark, G. A., "Failure Detection for Mechanical Structures: A State-Space Model Base Approach," Lawrence Livermore National Laboratory Report, in preparation.

2. Candy, J. V., and D. L. Lager (1979), "Identification, Detection and Validation of Vibrating Structures: A Signal Processing Approach," Lawrence Livermore National Laboratory, Livermore, California (UCID-18720).

3. Berman, A., and E. J. Nagy (1983), "Improvement of a Large Analytical Model Using Test Data," *AIAA Journal*, Vol. **21**, No. 8, August.

4. McCallen, D. B., and A. Astaneh-Asl (1988), "Computational Simulation of the Nonlinear Response of Suspension Bridges," *Proceedings of the World Structures Congress*, Elsevier, San Francisco, California.

5. McCallen, D. B., and A. Astaneh-Asl, "Explicit Dynamic Analysis of Cable Supported Bridges," in preparation, to be submitted to *Computers and Structures*.

6. McCallen, D. B., C. R. Noble, M. S. Hoehler, and M. Gerhard, "The Seismic Response of Concrete Arch Bridges with Focus on the Bixby Creek Bridge, Carmel, California," Lawrence Livermore National Laboratory Report to the California Department of Transportation, in preparation.

# Modeling and Simulation of Wireless Sensor Networks

**Rowland R. Johnson**
*Computer and Communications Engineering Division*
*Electronics Engineering*

Currently, Lawrence Livermore National Laboratory (LLNL) has several wireless sensor network (WSN) programs, and it is likely that the number will increase. In FY-98, LLNL researchers developed a core competency to model and simulate WSNs, which will give LLNL a competitive advantage in research and development. Contributions to existing programs are described, and the results of LLNL-funded academic research are presented.

## Introduction

Recent advances in micro-electromechanical systems, low-powered circuit technologies, and computer networking have made possible a low-powered WSN involving thousands or millions of sensor nodes. Over the next ten to twenty years, LLNL will have the opportunity and obligation to participate in programs involving these networks. Just as the capability to model and simulate nuclear explosions was an R&D enhancer for nuclear devices, the capability to model and simulate WSNs will be an R&D enhancer for WSNs. A modeling and simulation capability 1) will dramatically reduce the effort required to develop a WSN; 2) can be used as a planning tool for a deployment; and 3) can be used to evaluate proposed network systems.

Current wireless network techniques do not meet node power consumption and total system throughput requirements. Academic researchers are beginning to consider non-hierarchical network topologies as a means to address power consumption and system bandwidth issues. Data packets are routed from node to node based on which RF channel offers the most efficiency at a particular time (usually, this results in a data packet being routed to a nearby node). Higher RF channel efficiency allows a lower RF power level to be used, which reduces the node power consumption. It also reduces RF interference by other nodes, thereby increasing total system throughput. It has been shown that, although latency usually increases (due to increased hop count), power consumption decreases and system throughput increases.

The disadvantage of non-hierarchical network topologies is that their behavior is only beginning to be understood. In addition, research has been focused mainly on applying these techniques to personal communication systems, such as cellular telephones and personal digital assistants. It is unknown how effective these existing techniques will be for the typical LLNL WSN application.

## Progress

During FY-98 a core competency in modeling and simulation techniques for WSNs was developed. A team of LLNL researchers have attained expertise in analyzing WSNs via modeling and simulation techniques. Expertise has been developed in the use of state-of-the-art, commercially available simulation packages. In addition, researchers have developed key computer software components. Wireless sensor networks being considered are tightly coupled with a physical system that will affect their operation. Therefore, it is important to be able to integrate a simulator of such a physical system with the WSN simulation.

Consider, for example, a chemical attack warning system. Terrain and atmospheric conditions will affect both the movement and dispersion of a chemical release, and the ability of the wireless nodes to communicate with one another. An integrated simulation system makes it possible to investigate system issues such as the effect of terrain on battery life.

The development of the core competency has been guided by the requirements of two existing WSN programs at LLNL: the Joint Biological Remote

Early Warning System (JBREWS) and the Wireless Sensor project.

JBREWS will consist of 100 to 1000 biological sensor nodes, each capable of wireless communication with other sensor nodes or a central control node. The objective is to detect, track, and predict the movement of airborne biological agents. JBREWS is a multi-phase program, each phase based on newer, more sophisticated technology.

The performance of a JBREWS deployment will depend on sensor node placement, terrain, and atmospheric conditions. An overall system effectiveness is determined by obtaining performance criteria for widely varying deployment scenarios. Obtaining performance criteria for an actual deployment is expensive, and simulation offers a cost-effective alternative.

A primary concern of the JBREWS project is that the system be able to configure itself (**Fig. 1**) after sensor nodes have been randomly deployed, for example, by being air-dropped. In this case it is impossible to determine, *a priori*, what communication links between the nodes will be available. Once placed, each sensor node must be able to determine other sensor nodes that it can communicate with. Then, the sensor nodes, acting collectively, must configure themselves into a network. Several self-configuration algorithms were investigated.

JBREWS uses time-division multiple-access (TDMA) transport protocol as the basis for the communication channels. TDMA is characterized by a longer, though constant, latency when compared to other transport protocols. By eliminating the acknowledgment of successful data packet receipt, the latency will be reduced. The attendant reduction in reliability is dependent on the terrain and the exact placement of sensor nodes. Investigation of this issue by analytic means is not feasible.

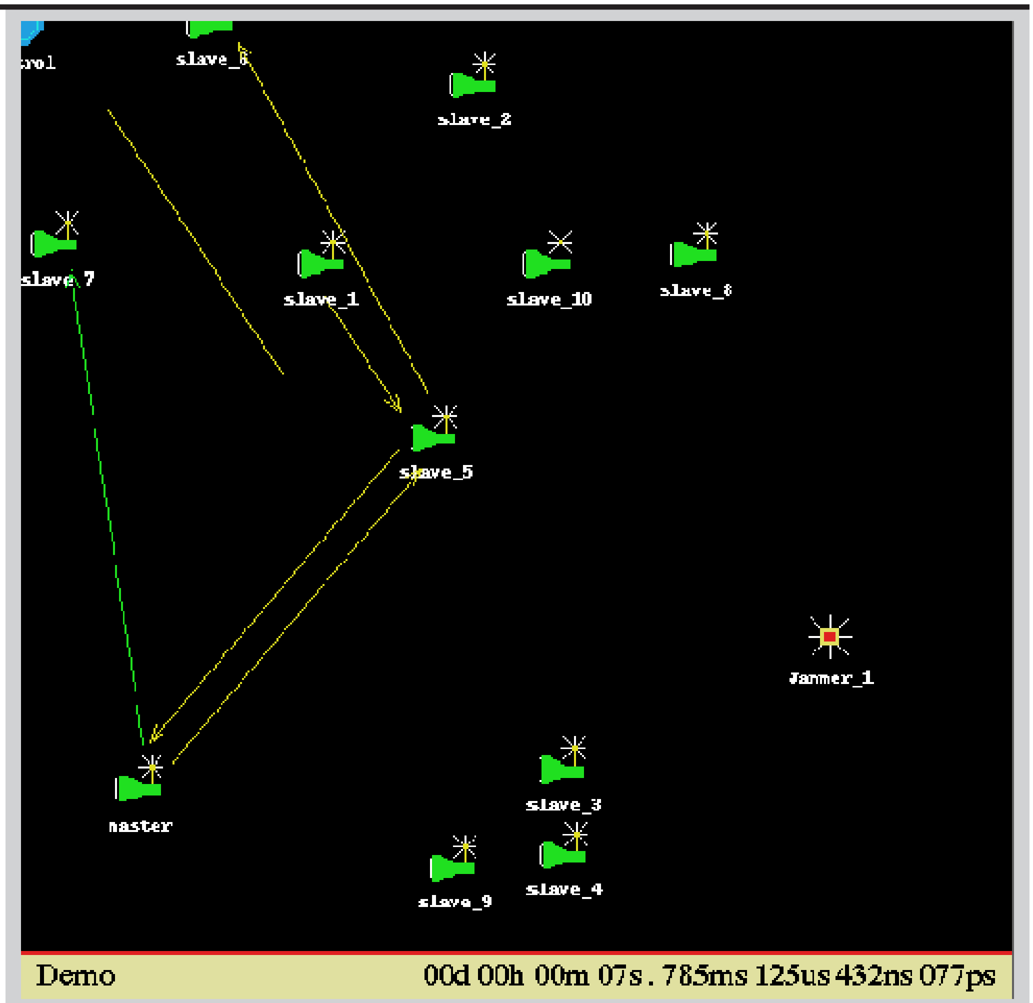


*Figure 1. The JBREWS wireless sensor network configuring itself. The master node has determined that it cannot communicate with slave_1 and slave_6 (because of the obstruction). Therefore, it relies on slave_5 to relay data packets to slave_1 and slave_6.*

However, simulation results have shown that the elimination of acknowledgment messages has had the unforeseen effect of reducing average latency, while increasing the variance in latency times.

The Wireless Sensor project is a joint project between LLNL's Microtechnology Center (MTC) and the University of California, Los Angeles (UCLA) to develop very low power wireless sensor nodes. The current focus of the project is on hardware development. In the next two or three years there will be opportunities to configure these nodes into systems for applications that are currently unknown. A simulation that accurately reflects current technology or technology being developed will provide the means to explore new applications.

The key issue that will drive the design of a system is the expected battery life (**Fig. 2**). This has a direct impact on system timing, because accurate clocks use more power than less accurate clocks. In turn, clock drift can degrade overall system performance by causing packet collisions and/or high latency times. The current simulation accurately reflects 1) the power cost of computing and communication operations, and 2) clock drift and system timing considerations.

In addition to work at LLNL, research at the University of California, Berkeley (UCB) on efficient multi-cast protocols was funded in FY-98. In many WSN applications there is a control node that will send directives to all sensor nodes. Typically, receipt of multi-cast packets is not acknowledged due to the network congestion that would occur when all sensor nodes acknowledge receipt at the same time. Without acknowledgment, totally reliable directive delivery is not possible, although it is possible to achieve some degree of reliability.

The research task undertaken by UCB was to develop methods to achieve high reliability with

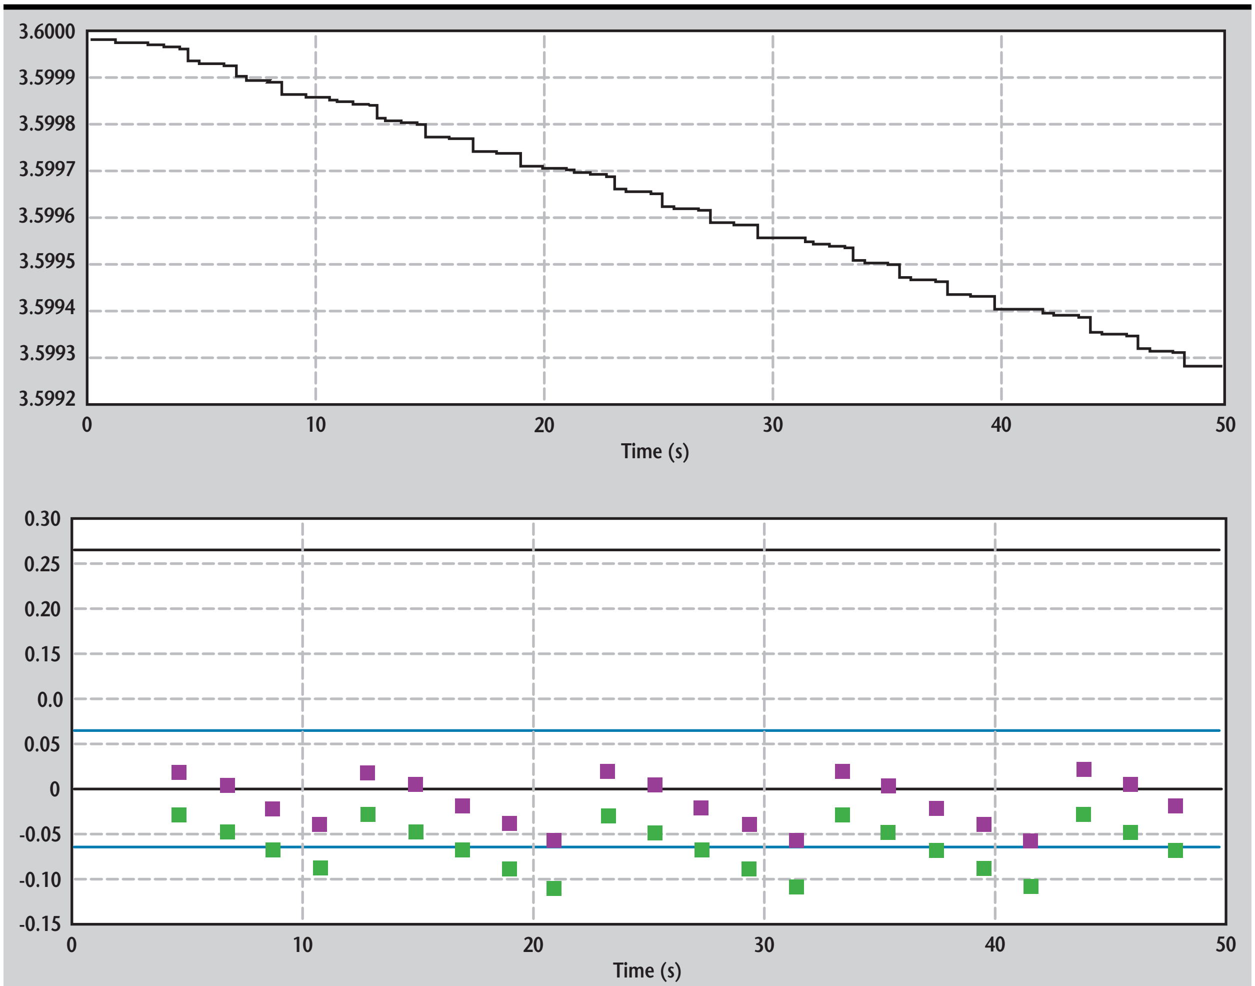


*Figure 2. Battery level and packet arrival times. (a) The battery is being drained whenever the receiver is turned on or the transmitter is used to transmit a packet; (b) shows a window in time when the receiver is turned on, to acquire packets being sent to it. To save power the receiver is turned off outside this window. A clock drift is shown by the fact that packets arrive earlier and earlier until finally a packet arrival falls outside the window. At this point data packets are being lost until the system can re-sync itself.*

non-acknowledged multi-cast protocols. Several multi-cast protocols were considered with the effectiveness of each being determined by modeling and simulation.

The result of this research is a "tunable" multi-cast protocol, that is, by specifying the degree of reliability, protocol parameters are selected that yield the desired reliability with the least overhead.[1,2]

## Future Work

Thus far, we have used commercially available simulation systems for the basis of our work. The disadvantage of this approach is that the simulation kernel can not be modified for the purpose of 1) achieving tighter integration with existing simulations of physical systems; or 2) porting the simulation kernel to a massively parallel computer. In FY-99 we will develop the capability to construct simulation kernels that are specific to a particular application.

## References

1. Wong, T., T. Henderson, S. Raman, and R. Katz (1998), "Exploiting Applications Level Framing in Reliable Multicast for Periodic Information Dissemination," Technical Final Report, University of California, Berkeley, August.

2. Wong, T., T. Henderson, S. Raman, A. Costello, and R. Katz (1998), "Policy-Based Tunable Reliable Multicast for Periodic Information Dissemination," Technical Final Report, University of California, Berkeley.

# Information Warfare Analysis Capability

**J. C. Smart**
*Electronics Engineering Technologies Division*
*Electronics Engineering*

The Information Operations, Warfare, and Assurance initiative at Lawrence Livermore National Laboratory has advanced the enabling core technologies for information operations analysis. Special emphasis was placed on computer networks and telecommunications systems.

## Introduction

With the rapid growth of global computing and communications, information security is a critical issue in all discussions of protection of the national infrastructure. The purpose of our project—the Information Operations, Warfare, and Assurance (IOWA) initiative—is to advance the enabling core technologies of this field.

Special emphasis was placed on computer networks and telecommunications systems.

## Progress

During FY-98, we developed 1) techniques for identifying the topology of large, complex computer networks; 2) data representation models for these systems; 3) high-performance methods for visualizing the resulting complex models; 4) automated analysis methods for processing large-network representations; 5) specialized search techniques for isolating vulnerabilities; 6) a foundation for simulating network operation; and 7) an assessment methodology for determining the consequences of system component failure or disruption.

To automate information system protection, it is necessary to first identify the visible components that an intruder might attempt to access, and to determine the specific information that might be inferred about each component.

We began by developing a set of software modules for analyzing Internet-related protocols. This software examines the information that flows across a computer network and extracts network topology and details about the configuration of each component. At present, the tool suite processes over 20 popular Internet protocols, retrieving over 50 different system operating parameters.

Since modern computer networks have grown considerably in size (that is, more than 25,000 nodes), a special model was developed to capture the enormous amount of information that the tools process. The resulting model uses a unique blend of relational database technology integrated into a graphical theoretical framework, providing rapid information retrieval in an environment conducive to large-network mapping and analysis.

We demonstrated a platform-independent viewer for browsing the graphics model with integrated access to the relation database. To better manage the complexity of large networks, several powerful dependency constructs, graphics operations, and reduction functions were incorporated into the model. A diverse suite of generic graph-, fault-tree-, and Internet-specific processing algorithms was developed and demonstrated.

To better understand the nature of computer and network vulnerabilities, a taxonomy of known vulnerabilities was developed that forms the basis of our new vulnerability database. This database was subsequently populated with vulnerability facts from industry and private sources. The end result is a tool that can now be used to automate the search for weaknesses in our computer systems.

Working with the modeling tools, an environment was constructed to perform high-fidelity simulations of computer networks. The resulting tools can be used to simulate computer networks captured in the IOWA model. Arbitrary computer networks can also be constructed in the simulation environment and used to generate network traffic to test and calibrate the network mapping tools.